LA-UR-13-22126
Approved for public release; distribution is unlimited.

Title:	Plasma Paradigm Shift Document for the CoCoMANS Project
Author(s):	Knoll, Dana A
Intended for:	Report

Plasma Paradigm Shift Document for the CoCoMANS Project

D. A. Knoll, L. Chacon, W. Daughton, J. Payne

March 25, 2013

1 Introduction

This is a dynamic document in the Computational Co-design of Multi-scale Applications in the Natural Sciences (CoCoMANS) LANL LDRD DR project. As of the current date we are approximately half way through this three year project and we have made significant prototype-app and mini-app progress in our plasma application. We are progressing towards developing compact-app software and using this software to demonstrate a paradigm shift in computational plasma physics via an evolving computational co-design process.

First we must define a *paradigm shift demonstration* within the CoCoMANS project. We consider there to be three aspects: 1) an "effective" use of hierarchical parallelism on emerging architectures, 2) Significant new physics simulation capability not achievable with traditional approaches via enhanced model fidelity, increased numerical accuracy, or both, and 3) running on multiple types of hardware (GPUs vs multicore) with no "physics / numerics code" changes. This third aspect will be accomplished with heavy use of C++ templating and virtual classes. All of this should be accomplished within a documented computational co-design process. In this document we will go into some level of detail mainly on the second point, with some detail on the first point. More details on the first point and details on the third point will be addressed in [5].

Within the plasma application, the CoCoMANS project will demonstrate a paradigm shift in multiscale kinetic plasma simulation. Specifically, a fully ionized collisionless plasma of hydrogen ions, free electrons, and electromagnetic fields (**E** and **B**). The computational challenge is to integrate an ion-electron kinetic system (plus Maxwell's equations) on an ion time-scale and a system length-scale while retaining electron kinetic effects accurately. Such problems are truly multi-scale, primarily driven by the disparity in ion and electron masses.

To appreciate the spread in relevant scales consider Figure 1. Typically, global scales L are bigger than ion kinetic scales ρ_i. Then electron space scales are $(m_i/m_e)^{1/2}$ smaller, and electron time scales are m_i/m_e faster. The fastest time-scales are plasma oscillations ω_{pe} and light waves, while shortest spatial scale is the Debye length λ_D. It is important to note that explicit PIC codes

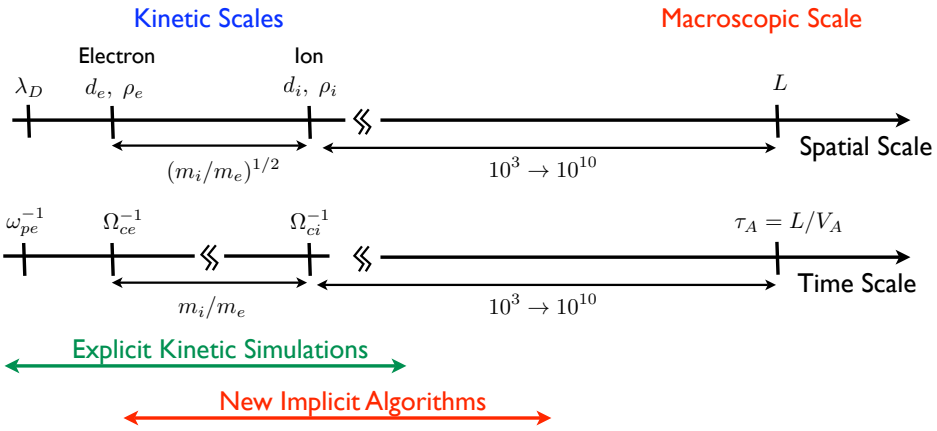

Figure 1: Kinetic Plasma Scales

(such as VPIC [1]) must resolve all time and space scales globally for stability [2]. Holding L/ρ_i fixed, computational cost (FLOPS) scale as $(m_i/m_e)^{5/2}$ in 3D. This steep scaling forces the user of an explicit algorithm to use "nonphysical" electron mass, usually $m_i/m_e = 25 - 100$ in large systems. Furthermore there are potential issues with energy conservation in explicit PIC codes [2]. This can render the accuracy of long time-scale (ion-time scale) simulations questionable.

In CoCoMANS we are developing a moment-based scale bridging algorithm which will simultaneously have improved algorithmic properties for long time simulation and provide for more efficient utilization of emerging architecture resources (computational co-design). The Underlying algorithm is a combination of a High-Order (HO) kinetic system and a Low-Order (LO) moment system which results from phase-space moments of the kinetic system.

The HO and LO problems work together to provide algorithmic acceleration to the HO solver. The LO PDE problem solution provides the electromagnetic fields ($\mathbf{E}$ and $\mathbf{B}$) to the HO problem. Closures required to solve LO problem (ion and electron stress tensors) come from the HO problem solution (ion and electron distribution function). In a "multigrid sense", the LO problem is a coarse space correction for the HO problem. The KEY algorithmic aspect is that an appropriately defined LO problem can advance the EM fields using a coarser mesh and larger time step as compared to the brute-force explicit PIC algorithm, while conserving energy. We refer to this evolving algorithmic idea as implicit PIC. However, it is important to understand that the HO solver is sub-cycled in order to follow complex particle orbits accurately. This algorithm is substantively different from an explicit PIC algorithm. Algorithmic details of these ideas as can be found in [3, 4].

It is useful to contrast this approach with the traditional explicit approach used in VPIC [1]. VPIC has been an exceptionally successful plasma simulation code, with important first-of-a-kind simulations on a number of important plasma physics problems. It uses a traditional explicit algorithm and thus

its characteristics are optimal for electron time-scales and Debye length space-scales. VPIC has been made very efficient on a number of modern computing platforms. It has achieved approximately 25% efficiency of the peak floating point potential of some architectures, since explicit PIC codes are limited by the available memory bandwidth. This efficiency is decreasing on emerging architectures. Since VPIC uses an explicit algorithm it does not conserve energy. This is typically not a large issue for simulating problems on electron time-scales. However, this can become an issue for long time (many global Alfven times) simulations done with realistic ion to electron mass ratios. Additionally, it should be pointed out that in an explicit algorithm the field solve is executed once per particle time step. In some explicit codes (not VPIC) this field solve will require global communication.

Again, more detailed discussion of out prototype/ algorithm progress can be found in [3, 4] and documentation of our mini-app progress and our parallelization strategy can be found in [5]. Here we focus on two physics problems we will use to establish a paradigm shift in computational plasma physics, and what we plan to demonstrate in each case. The first problem, the slow shock, is 1-D, 3-V, electromagnetic. It will provide the earliest insight into the impact of the new algorithm and it efficient implementation on multi-node, many-core + GPU, hardware. The second problem, island coalescence, is 2-D, 3-V electromagnetic. This problem will be significantly more taxing in terms of required computational resources, and will be our trophy example near the end of the project.

In the CoCoMANS project, we are demonstrating that the combination of advanced algorithmic development, along with the efficient use of advanced architectures, will provide a more substantial scale-bridging capability, as compared against solely relying on the hardware. When using a traditional explicit algorithm one is counting solely on the hardware when bridging from electron time and space scales up to ion time scales and system space scales. The CoCoMANS philosophy is: *The development of new hierarchical algorithms within a computational co-design process is key to achieving effective utilization of emerging hierarchical architectures is* .

2 Slow Shock Problem

Collisionless shocks are an attractive electromagnetic test problem since they can be studied in 1D and many of their basic properties are well-established in the literature. To narrow our focus, we will consider slow-mode shocks associated with the slow magnetosonic wave [6]. The present knowledge is based on two-fluid theory [7] together with hybrid (kinetic ions and fluid electrons) simulations [8, 9, 10, 11] and fully kinetic simulations (both implicit [12] and explicit [13, 14, 15]). For sufficiently oblique shocks, there is evidence that the electrons are heated in the parallel direction much more strongly than observed in hybrid simulations. The leading explanation is that backstreaming ions drive an obliquely propagating kinetic Alfven wave which is Landau resonant with

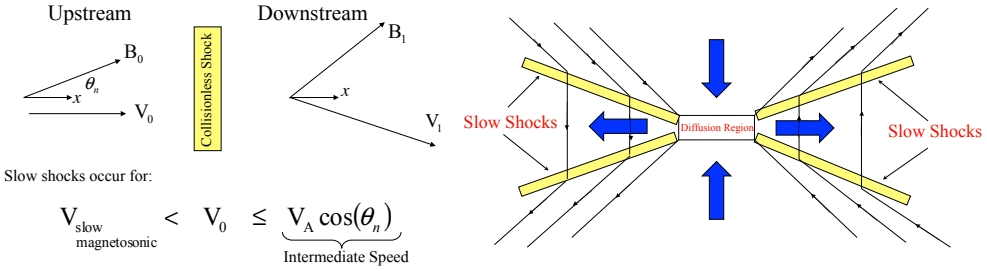

Figure 2: Left panel illustrates the shock notation in the normal incidence frame (i.e. shock frame) while the right panel illustrates the standing slow-mode switch-off shocks which are thought to occur for the Petschek model of fast reconnection.

electrons [13, 14, 15]. Thus it appears that overall structure of the oblique slow shocks is set predominantly by ions, but there are still some interesting electron kinetic effects to explain the heating and dissipation.

This physics is potentially quite interesting, but is very difficult to study with explicit simulations, which require millions of time steps even with artificial ion to electron mass ratio $m_i/m_e \sim 25$. These explicit simulations eventually encounter significant numerical heating which can make the results difficult to interpret. To make further progress and consider larger systems at more realistic mass ratio, this problem really needs implicit energy conserving algorithms.

With this goal in mind, in the following sections we provide some details for setting up kinetic simulations of slow-mode shocks, and briefly discuss what we expect to demonstrate.

2.1 Terminology and Basic Setup

The left panel of Fig. 2 illustrates the notation in the normal incidence shock frame (i.e. the shock is stationary and inflow is normal to the plane). The upstream magnetic field and velocity are given by $\mathbf{B}_o$ and V_o while the corresponding downstream values are given by $\mathbf{B}_1$ and V_1, and the angle between the shock normal and the upstream magnetic field is denoted by θ_n. Slow-mode shocks occur when the upstream flow velocity V_o is greater than the slow magnetosonic speed but less than or equal to the intermediate speed. Going from upstream to downstream, both the density and velocity *increase* while the magnetic field *decreases*. The limit $V_o = V_A \cos(\theta_n)$ corresponds to a *switch-off* slow mode shock since the transverse component of the magnetic field vanishes in the downstream $B_{1y} = B_{1z} = 0$. The majority of simulation studies in the published literature have focused on the switch-off limit, since standing slow-mode switch-off shocks are an essential feature in the Petschek model of fast reconnection as illustrated in the right panel of Fig. 2.

The most common way of setting up shocks within kinetic simulations is the *piston method* in which a flow is reflected off a fixed wall (or piston). Starting

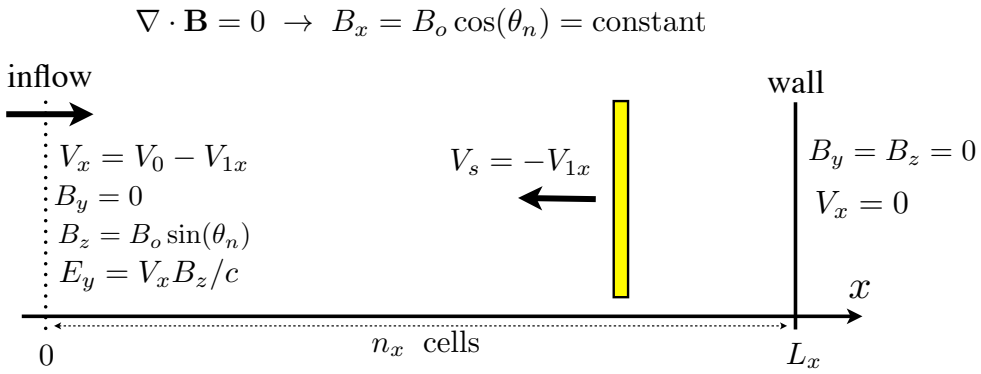

Figure 3: Kinetic simulation setup in the frame of the wall.

with the normal incidence frame, we perform an additional coordinate transformation such that $V_x = 0$ in the downstream region. This corresponds to performing the simulation in the frame of the piston, which is accomplished by solving the Rankine-Hugoniot relations assuming a Maxwellian plasma on both sides of the shock as described in Ref. [6] (see pages 11-13). Thus for a given set of upstream plasma conditions and shock angle, we can compute the expected down stream velocity V_{1x}. Switching into a frame where this is stationary, the shock will propagate to the left in the simulation frame illustrated in Fig. 3.

2.2 What we expect to show

- We will demonstrate energy conservation out to long times. This will allow a significant level of confidence in interpreting the numerical results to phenomena in the model, and not to numerical integration error.
- We will use a realistic mass ratio (1836), and test the sensitivity of our numerical results as a function of artificial mass ratios (25 and 200).
- We will be able to use a coarser spatial mesh with the advanced algorithm which will result in fewer particles and less memory requirements (another big plus on emerging architectures). Since this physics problem has no charge separation effects, our new algorithm can use mesh spacings which are large compared to Debye lengths, yet fine enough to resolve all important physical gradient scale lengths. This savings could be very significant.
- We expect efficient turn around by using on order 20-30 nodes on moonlight. This is a result of our advanced computational co-designed algorithm which will have maximum on-node computation for limited node-to-node communication. We will provide detailed profiling of these simulations to demonstrate a large ratio of computation to communication.

We expect to be able to produce new result providing physics insight into this problem by running out to longer times with real mass ratio and energy conservation. We will produce at least one physics publication describing these new results. This study will be completed in FY-13, and serve as an initial demonstration of a paradigm shift. This effort should prepare us well for a more complex and demanding 2-D, 3-V paradigm shift demonstration in FY-14 using the island coalescence problem.

3 Island Coalescence Problem

For the final demonstration of a paradigm shift in computational plasma physics we will consider a more computational demanding 2-D, 3-V problem. We consider a historical, and well studied, magnetic reconnection problem, island coalescence [18, 19]. Our team has simulated this problem using resistive MHD [20], Hall MHD [21], and VPIC [22]. The physics of this problem is an intimate interaction between the system scale, current sheet space-scales, the ion inertial space-scale (d_i) and non-Maxwellian electron effects. The range in spatial scales extends from the system scale (a magnetic island width) down to the ion inertial length, and down to the electron inertial length when reconnection is triggered.

In this problem two magnetic islands are attracted to one another by standard MHD forces. As the islands approach each other a current sheet is formed where the magnetic flux is compressed between the two islands. A resistive MHD fluid picture of this process can be seen in Fig. 4. Here we can see a current sheet forming as the two islands come together. For large enough resistivity in a collisional fluid model, the two islands will directly merge (coalesce) into one island via classic magnetic reconnection. However, if resistivity is low enough then the islands may bounce or slosh as a result of build up of magnetic pressure in the current sheet [20, 21]. In extended fluid models, such as Hall MHD, if the current sheet is compressed down to a scale of d_i prior to sloshing then magnetic reconnection will proceed independent of the magnitude of resistivity [21].

Fluid simulations which generate current sheets on the scale of the ion inertial length most likely will be missing kinetic effects. However, system scale kinetic simulations of this problem are extremely challenging. A recent, very impressive, 2-D VPIC simulation of this problem [22] used 20,480 cores on the Pleaides cluster (NASA), required a spatial grid of 17920 x 8690 cells, and require 600,000 time steps. This simulation used an artificial ion / electron mass ratio of 25 in order to render the simulation out to ion time-scales tractable. Again, an explicit PIC algorithm, such as the one in VPIC, must maintain all electron time and space scales. Thus VPIC has restrictive constraints on grid spacing and time step size. This is a brute force kinetic simulation which must depend on primarily on advanced computing hardware for acceleration. A snapshot in time from the VPIC kinetic simulation can be seen in Fig. 5. Here we can see non-Maxwellian (kinetic) effects can evolve in the current sheet.

The collisionless kinetic results in [22] have shown results which are different

from fluid models [20, 21] for cases with island widths on order 50 - 100 d_i. For this regime, the kinetic solution does not appear to slosh after the first initial bounce. For a fluid simulation, in this same situation, we observe sloshing. It is natural to expect difference between fluid simulations and collisionless kinetic simulations. For instance, fluid simulation will require a finite resistivity which will always produce some level of magnetic reconnection. The fluid simulations in [21] are slight smaller than the "large island" simulation in [22], and perhaps the resisitivity is to large for an accurate comparison with collisionless kinetic simulations. Thus we may need to perform a few more fluid simulations for a more representative comparison. (Note that if we drop the resistivity too far we may get into a multi-plasmoid instability regime)

In demonstrating our paradigm shift, a CoCoMANS goal is to first understand if the "large island simulation" in [22] is robust to both energy conservation, and a more realistic mass ratio. We can test this with our 2-D compact app software on advanced multi-node, many-core + GPU, hardware. If the "large island simulation" in [22] is indeed robust to these issues then we will work to elucidate the underlying physics behind the difference in fluid and kinetic simulations. As pointed out in [22], these differences could be driven by difference in observed ion and electron outflow speeds from the reconnection region as compared to fluid simulations. This would indicate stress tensors in the kinetic simulation are different from approximate closures used in the fluid simulation.

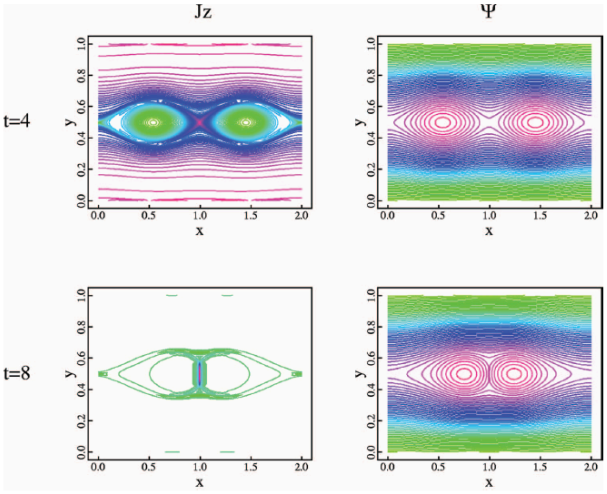

Figure 4: Fluid IC Results

Again, we emphasize the key points which will allow us to simulate the "large island problem" efficiently and accurately:

- Synergy between new algorithms and advanced hardware: Computational Co-design
- The implicit scale-bridging algorithm will be used to skip over fastest

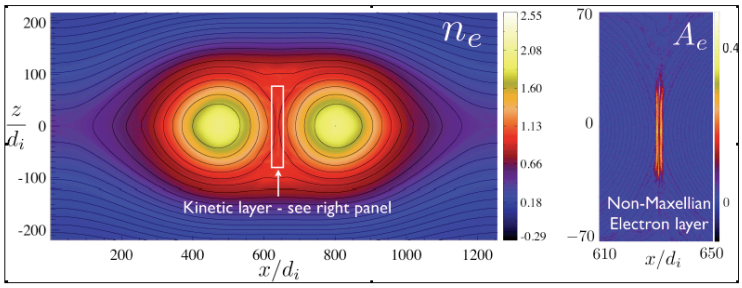

Figure 1. Kinetic VPIC simulations of island coalescence are now possible for island sizes approaching MHD scales. With initial island width w=100d_i this example required a spatial grid of 17920x8690 cells and was performed using 20,480 computational cores on Pleaides (NASA cluster). The left panel shows the electron density and flux surfaces (black lines). The right panel is a close-up of the region indicated showing the strongly non-Maxwellian character of the electron distribution as measure by the electron agyotropy $A_e = |P_{\perp 1} - P_{\perp 2}| / |P_{\perp 1} + P_{\perp 2}|$. These non-Maxwellian deviations play a crucial role in breaking the frozen-in condition to permit reconnection.

Figure 5: VPIC IC Results

LO problem time-scales: plasma frequency, etc - which are not directly relevant to our application. Energy conserving implicit methods will allow cell sizes $\Delta x \gg \lambda_D$.

- The HO problem will be sub-cycled to accurately follow electron trajectories, while field evolution is followed on slower time scale using orbit-averaged moments as sources. This sub-cycling will allow for significant isolation of on-node work with minimal node to node communication. This is the big picture goal, increased flops and reduced data movement. Also, following rapid electron motion is well suited for GPUs.

- As compared to traditional explicit PIC algorithms, this approach will have fewer cells (factor of 400), thus fewer particle, and fewer field solve time steps (factor of 10).

- As compared to traditional explicit PIC algorithms, the total number of steps per particle will typically be larger as a result of iteration within a time step between the HO and LO problems and iteration within the HO problem on the implicit particle steps. This iteration provides exact energy conservation, and is done in a manner that effectively uses FLOPS with almost no data motion.

The CoCoMANS project has been fortunate to hire one senior staff member, one post-doc, and one post-masters who have brought with them significant algorithm and mini-app experience in the plasma application area. From this CoCoMANS has inherited one "early co-design iteration" with mini-app experience for standard PIC and implicit PIC on CPU-GPU hardware [23, 24]. Furthermore, CoCoMANS has benefited from the LANL co-design summer school

in 2012 where a stand alone HO solver on the island coalescence problems was used as a model problem with both multi-core and GPU implementations [25]. This, along with the VPIC results in [22], have allowed us to make some estimates for the computational requirements for the "large island problem" on some available hardware.

Below are our current estimates for VPIC-equivalent 2D "large" island-coalescence simulation.

- Total number of electron substeps (3 Alfven times): $N_{\Delta t} = 3 \times 10^4 \times \frac{m_i}{m_e}$
- Domain of 500×1000 (vs. 10000×20000 on original VPIC run)
  - Mesh reduction allowed by implicit PIC.
  - Physics fidelity requires mesh packing for 2D island coalescence.
- 500 particles/cell (comparable to VPIC) $\Rightarrow$ $N_p = 2.5 \times 10^8$
- **Critical assumption**: electrons are pushed at rate of 2 ns/substep on GPUs. Basis:
  - 1D electrostatic application: 0.3 ns/substep [23]
  - 3D electromagnetic: 2 ns/substep [24] (with slower GPUs, but no cell-crossing treatment)

$$\text{Wall clock (days)} \sim \frac{0.13}{N_{GPU}} \times \frac{m_i}{m_e} \times N_{FE} \; ; \; N_{FE} = 1 - 10$$

m_i/m_e	$N_{GPU} = 80$ (30% Darwin)	$N_{GPU} = 960$ (10% Titan)
25 (as in VPIC)	1–10 hours	4-40 min
1800 (realistic)	3–30 days	6 hours - 2.5 days

This analysis is based on our partial replication model for parallelization [5]. Here, the geometry is replicated on each node and each node is given a share of the total number of particles. Nodes are required to communicate with each other only on the LO solver time step, Δt_{LO}. The nodes individually are able to sub-cycle the HO solver using time steps of Δt_{HO}, and this work is done in isolation. Typically we expect $\frac{\Delta t_{LO}}{\Delta t_{HO}} \approx 100$ for electrons. This along with the understanding that the HO solver will comprise more than 95% of our required computation effort will result in a large positive ratio between computation and communication.

4 Comparison with Traditional Implicit PIC

In addition to the above paradigm shift demonstrations, which are each focused around a particular open plasma physics questions, CoCoMANS will also execute a study which directly asses numerical accuracy of the new implicit

scale-bridging approach [3, 4] with the closely related original implicit moment method [12]. The results of this study will be documented and published. As stated many times in this document, the new scale-bridging algorithm conserves energy and charge, sub-cycles the HO problem to resolve particle orbits, and is demonstrably second-order accurate in time. These characteristics render the new approach [3, 4] far superior to it's predecessor [12] in terms of accuracy. This fact will be demonstrated and published. .

References

[1] Bowers, K.J. et. al., Phys. Plasmas 15, 055703 (2008).

[2] C.K. Birdsall and A.B. Langdon, in *Plasma Physics via Computer Simulation*, edited by W.W. Laing (Insitute of Physics Publishing, 1991).

[3] G. Chen, L. Chacon, D. Barnes, J. Comput. Phys. **230**, 7018-7036, (2011)

[4] W. Taitano, D. Knoll, L. Chacon, G. Chen, SIAM J. Sci. Comput. (2013, accepted)

[5] J. Payne, et. al. LA-UR-??? (2013)

[6] D.A. Tidman and N.A. Krall,*Shock Waves in Collisionless Plasmas*, Wiley-Interscience, (1971).

[7] F. Coroniti, *Laminar wave-train structure of collisionless magnetic slow shocks*, Nucl. Fusion bf 11,261, 1971

[8] D. Swift, *On the structure of the magnetic slow switch off shock*, J. Geophys. Res. **88**, 5685, 1983

[9] D. Winske, E. Stover and S.P. Gary, *The structure and evolution of slow mode shocks*, Geophys. Res. Lett. **12**, 295, 1985

[10] N. Omidi and D. Winske, *Structure of slow magnetosonic shocks in low beta plasmas*, Geophys. Res. Lett. **16**, 907, 1989

[11] N. Omidi and D. Winske, *Kinetic structure of slow shocks: Effects of the electromagnetic ion/ion cyclotron instability*, J. Geophys. Res. **97**, 14801, 1992

[12] H. Vu and J. Brackbill, *CELEST1D: An implicit fully kinetic model for low-frequency electromagnetic plasmas simulations*, Comput. Phys. Commun. **69**, 253-276, 1992

[13] W. Daughton, D. Winske, L. Yin and S.P. Gary, Role of electron physics in slow mode shocks, J. Geophys. Res. **106**, 25031, 2001

[14] L. Yin, D. Winske, W. Daughton and F.V. Coroniti *Dissipation in oblique slow shocks*, J. Geophys. Res. **110**, A09217, 2005

[15] L. Yin, D. Winske, W. Daughton and K.J. Bowers *Kinetic Alfven waves and electron physics: 2. Olblique slow shocks*, Phys. Plasmas **14**, 062105, 2007

[16] R. Morse and C. Nielson, Phys. Fluids **14**, 830 (1971).

[17] D. Forslund, Fundamentals of Plasma Simulation, *Space Science Review* **42**, p. 3-16, 1985

[18] D. Biskamp and H. Welter, Phys. Rev. Lett., **44**, 1069, (1980)

[19] P.L. Pritchett, Phys. Plasmas, **14**, 052102, (2007)

[20] D.A. Knoll and L. Chacon, *Coalescence of magnetic islands, sloshing, and the pressure problem*, Phys. Plasmas, **13**, 032307, (2006)

[21] D.A. Knoll and L. Chacon, *Coalescence of magnetic islands in the low resistivity, Hall-MHD regime*, Phys. Rev. Lett., **96**, 135001/1-4, (2006)

[22] H. Karimabadi, et. al., *Flux Pileup in Collisionless magnetic reconnection: bursty interaction of large flux ropes*, Phys. Rev. Lett., **107**, 025002 (2011).

[23] G. Chen, L. Chacon, D. Barnes, J. Comput. Phys., **231**, 5374-5388, (2012)

[24] Payne MS thesis, Nuc E MIT (2012)

[25] LANL Computational Codesign Summer School Report, 2012